à la carte: A Los Alamos Computer Architecture Toolkit for Extreme-Scale Architecture Simulation

Current approaches to building larger supercomputers—connecting commercially available SMPs with a network—may be reaching practical limits.

In response, the DOE Advanced Architecture Initiative seeks to research alternative high-performance computing architectures.

The path to better hardware design and lower development costs involves performance evaluation, analysis, and modeling of parallel applications and architectures, and, in particular, predictive capability.

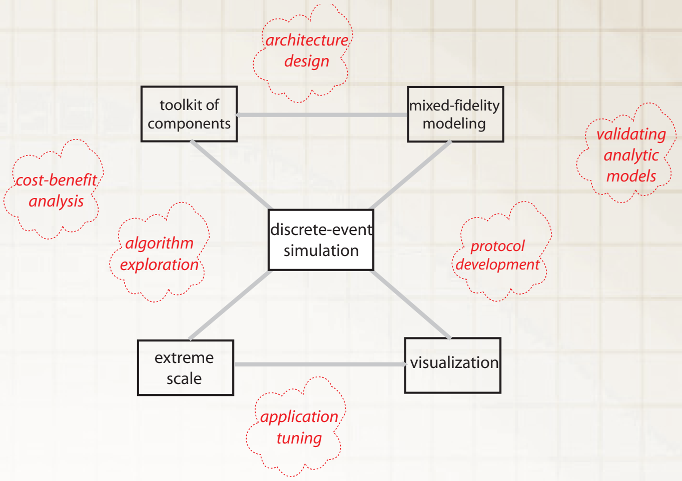

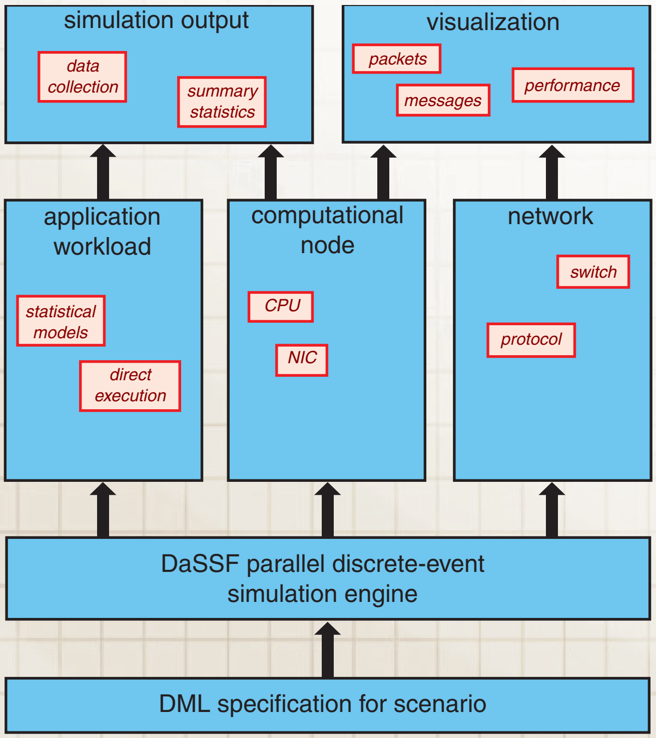

Approach

Our component-based design allows for the seamless assembly of architectures from representations of workloads, processors, network interfaces, switches, etc., with disparate resolutions, into an integrated simulation model.

Simulating systems of the size and complexity we envision requires efficient parallel simulation. We use a portable, conservative synchronization engine (DaSSF), developed by Dartmouth College, for the handling of discrete events.

Applications and computational workloads may be represented at a variety of fidelities:

- Simple random processes can load the hardware with message traffic having specified statistical properties.
- Direct-execution techniques allow one to run programs nearly exactly on real processors coupled to a simulated network.
- From time series of fine-grained simulations we will use neural networks and other learning algorithms to construct reduced models of the full system dynamics.

The initial prototype, comprising low-fidelity models of workload and network, models at least 4096 computational nodes in a fat-tree network. It supports studies of simulation performance and scaling rather than the properties of the simulated systems themselves.

Ongoing work is improving the fidelity of the representations and protocols.

Analysis

Our simulation provides complete detail concerning the history of messages and the propagation of packets through the simulated network.

We are currently running simulations of a 4096-node fat-tree containing 6144 switches using a simple random process to generate messages and a basic circuit-switched protocol.

Uniformly-Distributed Packet Destinations on 4096 SMPs

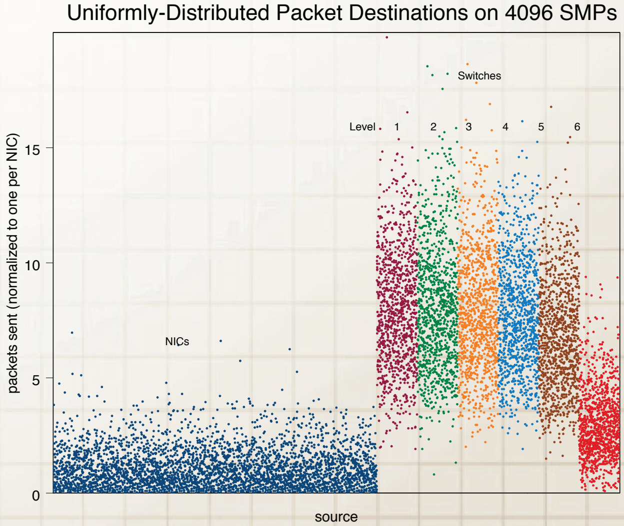

CFD-Distributed Packet Destinations on 4096 SMPs

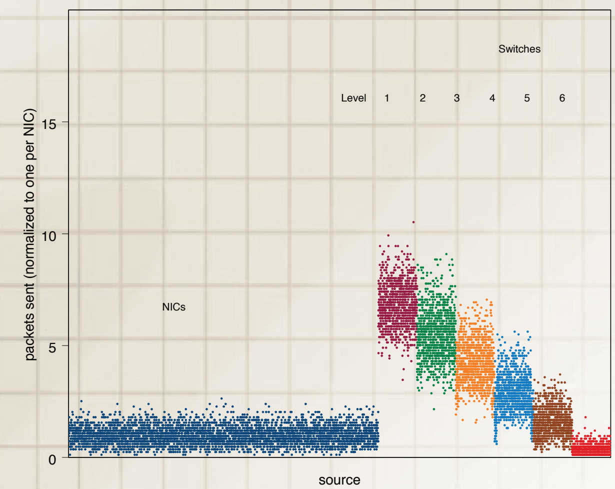

Scientific visualization techniques are used to analyze output as well as debug the simulations.

Visualization approaches include direct representations of the architecture as well as innovative abstractions of the architecture and dynamics of the system.

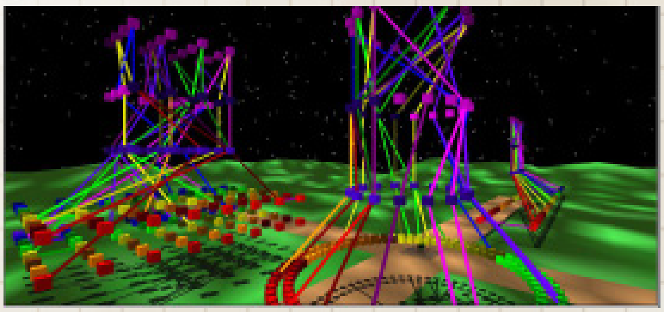

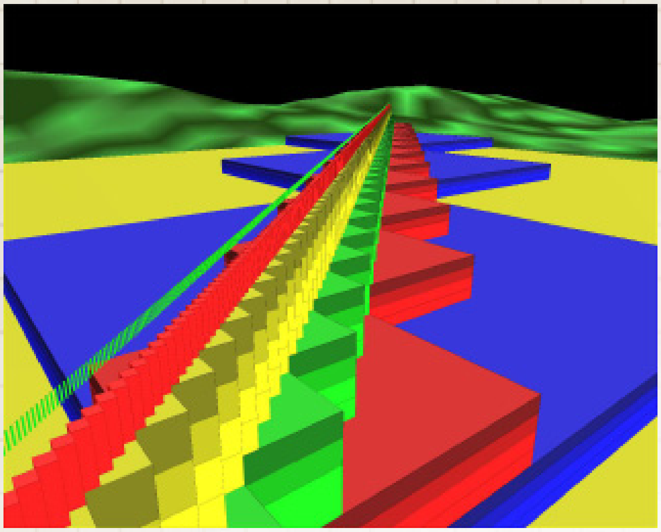

Contact Information

Send e-mail to parsim@lanl.gov.

Modeling, Algorithms and Informatics Group
Computer & Computational Science Division
Los Alamos National Laboratory
Los Alamos, NM 87545

November 2001

LALP-01-243

A US DEPARTMENT OF ENERGY LABORATORY

Los Alamos
NATIONAL LABORATORY
Los Alamos NM 87545